

# **An initial framework for the rapid qualification of long-term creep rupture strength via microstructural modeling**

---

**Applied Materials Division**

**About Argonne National Laboratory**

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see [www.anl.gov](http://www.anl.gov).

**DOCUMENT AVAILABILITY**

**Online Access:** U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (<http://www.osti.gov/>), a service of the U.S. Dept. of Energy's Office of Scientific and Technical Information

**Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):**

U.S. Department of Commerce  
National Technical Information Service  
5301 Shawnee Rd  
Alexandria, VA 22312  
**[www.ntis.gov](http://www.ntis.gov)**  
Phone: (800) 553-NTIS (6847) or (703) 605-6000  
Fax: (703) 605-6900  
Email: **[orders@ntis.gov](mailto:orders@ntis.gov)**

**Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI)**

U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831-0062  
**[www.osti.gov](http://www.osti.gov)**  
Phone: (865) 576-8401  
Fax: (865) 576-5728  
Email: **[reports@osti.gov](mailto:reports@osti.gov)**

**Disclaimer**

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

# **An initial framework for the rapid qualification of long-term creep rupture strength via microstructural modeling**

---

Applied Materials Division  
Argonne National Laboratory

September 2021

**Prepared by**

**A. Venkataraman, Argonne National Laboratory  
M. C. Messner, Argonne National Laboratory**



## **Abstract**

This report describes the development and testing of a new method for extrapolating short-term creep rupture test data to predict long-term rupture strength. The goal of this work is to reduce the time required to qualify new materials for nuclear service by reducing the lead time required for dedicated, long-term material testing to establish key long-term material properties. The new approach described here uses a physics-based model to predict the long-term creep rupture strength of 316H stainless steel using only short-term test data. The key idea is to use Bayesian inference to find the statistical distribution of the model parameters that best explain the short-term rupture data. Because the model is physics-based these parameters are all microstructural quantities that can be measured through detailed material characterization experiments. The Bayesian prior distributions provide a means for incorporating this characterization data into the final model to improve the accuracy of the long-term model predictions. However, where such data is not available the process still produces an accurate model based on an uniformed prior. Our hypothesis is that this approach more accurately extrapolates the short-term test data when compared to current, empirical methods. The report proves this hypothesis using actual long-term rupture data available for 316H, including tests with rupture times greater than 200,000 hours. The general approach developed here could be applied to other materials and other time-dependent material properties. Applying this new technique to develop long-term qualified material properties, potentially in conjunction with other accelerated qualification approaches like staggered qualification test programs, could greatly reduce the time required to qualify new materials for nuclear service.



## Table of Contents

Abstract	i
Table of Contents	iii
List of Figures	v
List of Tables	vii
1 Introduction	1
1.1 Background . . . . .	1
1.2 Organization . . . . .	6
2 Hu and Cocks model for 316H	7
2.1 Model description . . . . .	7
2.1.1 Slip rate . . . . .	7
2.1.2 Dislocation hardening . . . . .	8
2.1.3 Precipitation/solid solution hardening . . . . .	8
2.2 Implementation in NEML . . . . .	14
2.3 Verification . . . . .	15
2.4 Macroscale examples . . . . .	19
3 CPFEM model and throughput simulation database	21
3.1 Simulation setup . . . . .	21
3.2 Analysis . . . . .	23
4 Bayesian inference model and long-term predictions of creep rupture	25
4.1 Surrogate modeling . . . . .	25
4.2 Bayesian inference . . . . .	28
4.3 Assessing the results . . . . .	33
5 Conclusions and future work	37
5.1 Summary . . . . .	37
5.2 Future work . . . . .	38
Acknowledgements	39
Bibliography	41





## List of Figures

1.1	Example of traditional Larson-Miller correlation against creep rupture data. This example applies a quadratic polynomial regression linking the Larson-Miller parameter to the log of the rupture stress. This regression model, along with the associated prediction interval shown on the plot, can be used to calculate extrapolated design minimum stress to rupture data. . . . .	2
1.2	Example of three potential problems with empirical time-temperature (Larson-Miller) correlations. All three figures use a consistent labeling scheme where circles represent data, the line represents the model, black indicates time/temperature data within the experimental database used to calibrate the model, and gray indicates time/temperature points outside the model database (for example, points generated later after additional test time). . . . .	3
1.3	Flow chart demonstrating the framework described in this report for extrapolating short-term creep rupture data to predict long-term failure. . . . .	4
1.4	Illustration of a staggered test program providing increasing, leading qualified time-dependent properties. . . . .	5
2.1	Comparison between NEML implementation and [21] results for thermal aging at $T = 550^\circ \text{C}$ . . . . .	16
2.2	Comparison between NEML implementation and [21] results for thermal aging at $T = 600^\circ \text{C}$ . . . . .	17
2.3	Comparison between NEML implementation and [21] results for thermal aging at $T = 650^\circ \text{C}$ . . . . .	18
2.4	Unaged tension response of the Hu and Cocks model. . . . .	19
2.5	Aged tension response of the Hu and Cocks model. The solid lines are the unaged response, the dashed lines are the response after aging 100,000 hours at the test temperature. . . . .	19
2.6	Simulated creep tests using the Hu and Cocks model. . . . .	20
3.1	(a) The 100-grain microstructure used for all the simulations to generate the database of creep rupture times. (b) The loading conditions imposed on the microstructure. . . . .	22
3.2	Distribution of creep (a) strain and (b) strain rate curves from the different throughput simulations, compared with the experimental creep curve for the same loading conditions. These curves are all for a stress of 100 MPa. . . . .	22
3.3	Calculation of the increase in the excess volume (after elasticity) for (a) brittle creep failure and (b) ductile creep failure cases. . . . .	24
4.1	Summary histogram of the dataset after filtering. . . . .	26
4.2	Typical training history for fitting the surrogate model to the CPFE data. . . . .	27
4.3	Plot comparing the surrogate-predicted and actual CPFE rupture times, using the reserved test data. . . . .	28
4.4	Figure illustrating the concept of a hierarchical, pooled model to capture heat-to-heat variation in the model microstructural parameters. . . . .	30
4.5	Typical loss history for the inference optimization process. . . . .	31

4.6	Example normalized posterior distributions for the microstructural parameters inferred from the short-term rupture data. . . . .	32
4.7	Comparison between experimental data, the physics-based model, and a linear Larson-Miller model, plotting predicted rupture time as a function of the applied stress. . . . .	32
4.8	Relative error versus long term, $t > 100,000$ test data for the physics-based model with different values of $t_{limit}$ (i.e. the longest duration creep test kept in the inference database. . . . .	34

## List of Tables

2.1	Physical constants that do not vary heat-to-heat. . . . .	12
2.2	Physical properties that do not vary significantly heat-to-heat. . . . .	13
2.3	Physical parameters that will vary significantly heat-to-heat. Note that the precipitate number density would also need to be initiated, but could be done using Eq. 2.24 based on the values of radius and phase fraction given here. . . . .	14
2.4	Scaling factors for the internal variables. . . . .	15
3.1	The crystal plasticity model parameters varied in the throughput simulations to generate a database of creep rupture data. . . . .	21
4.1	Comparison of the accuracy of the physically-based Bayesian model as well as linear and quadratic Larson-Miller correlations in extrapolating long-term creep rupture data based on short-term data. For the Bayesian model, the table shows results from 10 repetitions of the complete process. . . . .	34



# 1 Introduction

## 1.1 Background

This report describes a new approach for qualifying long-term, time-dependent material properties for high temperature nuclear reactor structural materials. Specifically, the work described here focuses on extending the extrapolation window for creep rupture strength data to accelerate the qualification of new materials for nuclear service [1]. Creep rupture is a critical failure mode for future high temperature reactors [2], one that often limits the life of reactor structural components. By definition, creep rupture strength is a time-dependent property and requires long-term creep testing to establish reliable, qualified material data. Current material qualification approaches rely on a dedicated, long-term creep test program to produce data to fit an empirical model. Most codes and standards allow some limited extrapolation in time outside the direct test data using a time-temperature parameter [3–6] to trade higher temperature testing for shorter test times. However, as these approaches use empirical models to correlate and extrapolate the data they limit the allowable amount of extrapolation. For example, the ASME Boiler & Pressure Vessel Code Section III, Division 5 rules covering the design and construction of high temperature reactor structural components limits the allowed extrapolation to a factor of 3 to 5 in time, as described in Section III, Division 5, Subsection HB, Subpart B, Appendix HBB-Y. Many reactor designs call for 30 or even 60 year initial design lives in order to amortize plant capital costs. Given these targets, a minimum full-life qualification program would require creep rupture testing with times from 6 to 20 years before a new material could be entered into service.

While other factors also contribute to prevent rapid nuclear material qualification, this requirement for dedicated long-term test programs for critical time-dependent properties is currently one of the most unavoidable factors limiting the time between material discovery and qualification. Moreover, maintaining test programs of this type can be difficult, as it requires a substantial expenditure to maintain the program over a long period of time.

The current qualification approach relies on time-temperature parameters to trade temperature for time in extrapolating creep test results. Figure 1.1 illustrates the process using the Larson-Miller parameter applied by ASME [3]. The approach requires a large, long-term creep rupture database providing the test stress and temperature linked to the corresponding creep rupture time. This database should sample several commercial heats of the material, ASME requires at least three heats, so that the data is representative of a random future heat of material used in constructing the component. The approach correlates the Larson-Miller parameter

$$LMP = T (\log_{10} t_R + C) \quad (1.1)$$

with  $T$  the absolute temperature,  $t_R$  the creep rupture time, and  $C$  the Larson-Miller parameter, which unifies the effects of time and temperature to the log of the experimental stresses. The basic Larson-Miller approach selects a polynomial order based on the data and finds the polynomial interpolation of that order and value of  $C$  that best fits the data. Other approaches apply different time-temperature parameters, but the basic idea remains the same.

Figure 1.2 illustrates three potential problems with this basic version of time-temperature parameter correlation: overfitting, unreasonable extrapolation, and inability to capture a




Figure 1.1: Example of traditional Larson-Miller correlation against creep rupture data. This example applies a quadratic polynomial regression linking the Larson-Miller parameter to the log of the rupture stress. This regression model, along with the associated prediction interval shown on the plot, can be used to calculate extrapolated design minimum stress to rupture data.

mechanism shift. Of these three, overfitting is uncommon as typically low order polynomials are sufficient to match the data. However, the other two pathological conditions are common and are the primary reason why time extrapolation limits are required to ensure accurate long-term creep rupture predictions.

More sophisticated versions of the time-temperature parameter correlation approach include region split methods [7], which attempt to capture mechanism shifts using a piecewise polynomial correlation, with a breakpoint at the time/temperature location of the shift. This technique is only effective if the experimental conditions capture the mechanism shift — it does not work if the mechanism shift only occurs at times/temperature beyond the test conditions but potentially within the design life of the component. Moreover, the region split approach introduces additional user-selected hyperparameters, specifically the number and location of the breaks, which can potentially lead to overfitting.

This report describes a new approach to correlating and extrapolating creep rupture data. The goal is to extend the extrapolation window to reduce the number and duration of creep rupture tests required to qualify a material for a 30 or 60 year life. The new approach correlates creep rupture data using a physics-based, microstructural model for creep deformation and failure [8]. This model uses a Crystal Plasticity Finite Element (CPFE) approach to discretize the microstructure of the material and capture key aspects of the material structure. A single crystal model captures the details of grain bulk deformation, linking bulk creep to the underlying physical mechanisms of dislocation motion, precipitate hardening, solid solution strengthening, and bulk vacancy diffusion causing the deformation. The CPFE model includes a grain interface-cohesive model [9–13] to represent grain boundary sliding and diffusion assisted creep-cavitation and eventual grain boundary failure. The




Figure 1.2: Example of three potential problems with empirical time-temperature (Larson-Miller) correlations. All three figures use a consistent labeling scheme where circles represent data, the line represents the model, black indicates time/temperature data within the experimental database used to calibrate the model, and gray indicates time/temperature points outside the model database (for example, points generated later after additional test time).

authors and others have used similar models to accurately predict creep deformation and failure in a variety of materials [14–16]. Physics-based models can be more accurate in predicting long-term material behavior outside the test database. For example, models of this type can capture mechanism shifts induced by a transition from dislocation to diffusional creep in the grain bulk [17].

One key aspect of physics-based models is that the model parameters all have a physical meaning. For example, the grain boundary model parameters include the grain boundary diffusivity, the initial spacing of the cavity nucleation sites, related the grain boundary precipitate spacing, the grain boundary viscosity, etc. In theory, microstructural experiments and characterization could determine these parameters directly, without the need to correlate to creep deformation or rupture data. However, in practice building multiscale models of this type is difficult, particularly as the required microstructural data is often not available.

Instead, notionally, the approach developed here calibrates the microstructural parameters using Bayesian inference [18, 19] against short-term creep test data. This approach finds the statistical distribution of the model microstructural parameters that best explains the observed macroscale test data. Our hypothesis is that because the underlying model is physics-based the resulting calibrated model will accurately extrapolate to long-term conditions, even if only calibrated against short term data. *Specifically, we predict that a physics-based model of this type will extrapolate more accurately in time than current time-temperature parameter approaches, allowing for an increased time-extrapolation window and reducing therefore reducing the amount and time of long-term testing required to qualify a material.*

Moreover, as the model parameters are all notionally measurable microstructural quantities with definite meaning, the prior distributions used in the inference provide a natural way to incorporate any available microstructural information on the material. These informed priors help constrain the calibrated parameter distributions to remain physically-reasonable

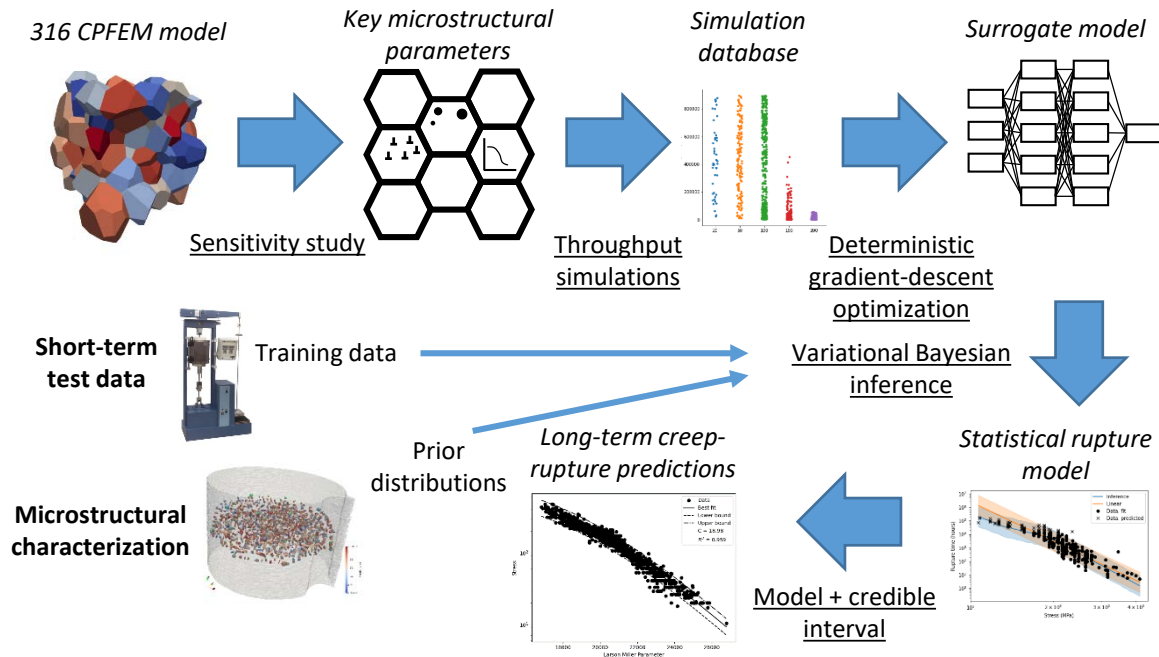


Figure 1.3: Flow chart demonstrating the framework described in this report for extrapolating short-term creep rupture data to predict long-term failure.

while still allowing the model to adapt to the available macroscale test data.

One complication to this approach is that the CPFE model is too numerically expensive to use directly in the inference process. Instead, we first generate a large database of simulation results, specifically creep curves and the associated rupture times, for grid of parameters including stress and the key model microstructural parameters as identified in a sensitivity study. Then, we fit a deep neural network surrogate model to this data, which maps stress and the key microstructural parameters to the resulting rupture time. This surrogate model is fast enough to use in the Bayesian inference step. As the surrogate model is quite accurate this extra step does not detract from the key idea of extrapolating using a physically-based model – basically the surrogate model perfectly reproduces the CPFE data. Figure 1.3 outlines the complete model development process used here to predict long-term creep rupture strength.

One additional advantage of the Bayesian inference approach is that the final model is statistical – it predicts not only the expected average rupture strength but also the expected distribution of strengths. This information is critical for developing design material data, which are often based on statistical lower-bounds of the expected heat-to-heat variation in the material properties.

Finally, this proposed approach is compatible with other strategies for accelerating material qualification. For example, one potential idea is a staggered qualification approach (see Fig. 1.4). Staggered qualification is compatible with the current restrictions on time extrapolation. Even with these factors, for short qualified design lives the required creep-rupture testing does not take long. For example, one year tests could support a five year qualified life. The staggered approach starts the required time-dependent testing, both long-



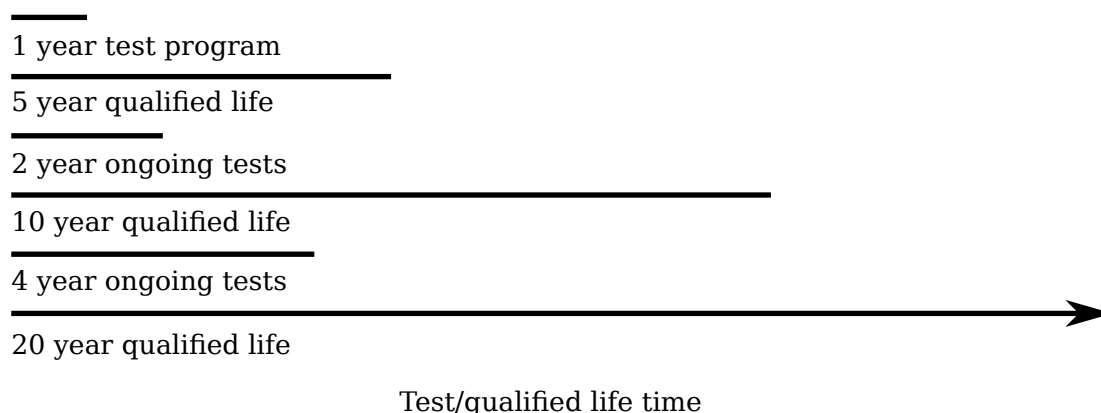


Figure 1.4: Illustration of a staggered test program providing increasing, leading qualified time-dependent properties.

and short-term, all at once. However, rather than waiting for all tests to finish to support a long 30 or 60 year qualified life instead the staged approach qualifies the material incrementally for longer and longer design lives as data comes in. So, for example, one year tests support an initial five year life. When the two year tests results are available the approach extends the qualified life to 10 years and so on. This approach would allow plant designers and operators to start operation based on the initial, short-term qualified properties. As long as the test program schedule provides new data before the old qualified life expires the staggered program would provide the data required for ongoing reactor operations, out to or beyond a 30 or 60 year target design life. However, if the plant designer wanted to eventually achieve these long target design lives, they would need to design the components initially with *expected* 30 or 60 year design material properties, not the initially-qualified short-term properties. These initial long-term predictions could be aggressively extrapolated, as they would not be the actual, qualified properties. If the later test data supports these initial properties (i.e. the material is at least as strong as initially predicted) components operating under such a staggered qualification program would be acceptable. The risk for the plant operator is that the later test data may not support the initial predictions, meaning the qualified component life would be shorter than expected, the component would need to be retrofitted, plant operations altered to reduce the load on the component, etc.

The physics-based method described here is compatible with a staggered approach to qualification as it could provide the initial, aggressively extrapolated design properties with better accuracy than a conventional time-temperature extrapolation. Introducing physics-based modeling to codes and standards at the same time as a staggered qualification method could ease the adoption of the new modeling technique, as the staggered test program would provide a conventional backstop for the simulation results. The ongoing accuracy of the long-term predictions against the test data would then also provide evidence of the effectiveness of the new modeling and simulation based approach described here, potentially meaning the next qualified material would not require ongoing testing.

While the current report focuses on creep rupture as a key material property for next generation high temperature reactors, the general approach detailed here could be applied to other time-dependent material properties like thermal aging, creep deformation, and even environmental effects like corrosion and radiation damage. The general strengths of the method — better accuracy in long-term extrapolation and a direct connection to the material microstructure — similarly apply for these types of material property predictions.

This report demonstrates the application of this new physics-based time extrapolation approach to predict the long-term creep rupture strength of 316H stainless steel. 316H is already a qualified ASME material for high temperature nuclear components, supported by an extensive long-term testing database extending beyond 200,000 hours [20]. We selected this material specifically because the ground truth creep data is available — this allows for a direct comparison of the long-term creep rupture strength predicted by the new approach based on a limited set of the rupture database (for example, limiting the inference process to rupture data with durations less than 15,000 hours), against actual long-term test results. This approach of demonstrating the potential predictive power of physics-based models against existing material test data is one way to build confidence in the new extrapolation method.

The results here demonstrate that the new, physics-based approach can in fact accurately extrapolate short term creep rupture data to long-term predictions. Moreover, we demonstrate that this new modeling approach is substantially more accurate than simple time-temperature parameter extrapolation from the same database.

## 1.2 Organization

Chapter 2 describes the implementation of a physically-based model for 316H developed by Hu and Cocks [14] into a MOOSE-based CPFE framework. This model is completely physically based in that it decomposes the strength opposing dislocation creep in the grain bulk into four mechanisms, each of which has a direct connection to some underlying microstructural property of the material. Our original plan was to use this model to represent the grain bulk deformation in 316H, combined with the existing grain boundary model. However, due to a delay in the start of the work we instead proceeded in parallel generating the throughput simulation database and surrogate model using a simpler, but still physically based creep model, already implemented in the MOOSE CPFE framework. Chapter 3 describes this model along with a description of the throughput simulation database. Chapter 4 then describes the process of fitting a surrogate model to the data and completing the model by calibrating property distributions using Bayesian inference against the time-limited 316H rupture database. This chapter demonstrates the ability of the final model to accurately predict long-term creep rupture and contrasts the accuracy of the new approach against conventional time-temperature extrapolation. Finally, Chapter 5 summarizes the work described here and details future work on applying this new approach to accelerate the qualification of new nuclear structural materials.

## 2 Hu and Cocks model for 316H

### 2.1 Model description

This chapter describes a single crystal, physics-based model for the creep deformation of 316H stainless steel developed by Hu and Cocks and details its implementation in a MOOSE-based CPFEM framework. Hu et al. (2020) [21] and Hu’s dissertation [22] describe the model in sufficient detail to reimplement the model. The model evolved over time as the authors added improvements. The version implemented here follows [21] with some changes described below.

The model represents dislocation creep in 316H as affected by three time-temperature dependent strengthening mechanisms:

1. Dislocation forest hardening
2. Precipitate strengthening caused by the Laves ( $\text{Fe}_2\text{Mo}$ ) and carbide ( $\text{Cr}_{23}\text{C}_6$ ) phases
3. Solid solution strengthening caused by Cr, C, and Mo.

In addition, the full model includes a backstress contribution to flow caused by dislocation pileups and the effect of solute drag on the material flow rate. Neither of these mechanisms are included in the version implemented here as both have a relatively small effect on the creep response of 316H.

We implement the single crystal model in NEML, a material model library developed at Argonne National Laboratory. NEML is in turn connected to MOOSE, where we execute the CPFEM modeling discussed in Chapter 3. NEML provides a framework for embedding a slip system model into general crystal plasticity kinematics. For a general description of the kinematic framework see [8, 23, 24] or the NEML documentation at <http://neml.readthedocs.io>. To implement the model then we only need to implement the slip-system level slip rate and hardening models.

The following provides a mathematical description of the model. Tables 2.1-2.3 define the parameter values.

#### 2.1.1 Slip rate

The slip rate equation is

$$\dot{\gamma}_i = \dot{\gamma}_0 \exp \left[ -\frac{\Delta F_0}{kT} \left( 1 - \left| \frac{\tau_i}{\tau_{CRSS,i}} \right|^c \right)^d \right] \text{sign}(\tau_i) \quad (2.1)$$

where  $\dot{\gamma}_i$  is the slip rate on system  $i$ ,  $\dot{\gamma}_0$  is the reference slip rate,

$$\Delta F_0 = \alpha_0 G_0 b^3, \quad (2.2)$$

with  $G_0$  the shear modulus at absolute zero,  $b$  the Burgers vector, and  $\alpha_0$ ,  $c$ , and  $d$  are parameters describing the energy barrier against climb-assisted dislocation glide,  $k$  is the Boltzmann constant,  $T$  the absolute temperature,  $\tau_i$  is the resolved shear on system  $i$ , and  $\tau_{CRSS,i}$  is the slip resistance.

The slip resistance on each system is the sum of three terms representing forest hardening, precipitation hardening, and solid solution strengthening

$$\tau_{CRSS,i} = \sqrt{\tau_{d,i}^2 + \tau_p^2 + \tau_s^2} \quad (2.3)$$

with  $\tau_{d,i}$  the forest hardening strength on system  $i$ ,  $\tau_p$  the precipitation hardening strength, and  $\tau_s$  the solid solution strength. The precipitation and solid solution strengthening terms are coupled, as the growth of precipitates pulls atomic species out of solution. These terms combine the hardening effects of the Laves and carbide phases. Each slip system has its own value of the forest hardening term but they all share the same precipitation and solid solution strengths.

### 2.1.2 Dislocation hardening

The dislocation forest hardening model is fairly standard, combining both self- and latent-forest hardening with thermally-activated recovery. The strength is given in terms of the dislocation spacing on the slip plane:

$$\tau_{d,i} = \frac{\alpha_d G b}{L_{d,i}} \quad (2.4)$$

with  $\alpha_d$  the interaction coefficient,  $G$  the temperature-dependent shear modulus, and  $L_{d,i}$  the dislocation spacing. The dislocation spacing evolves with the differential equation

$$\dot{L}_{d,i} = -L_{d,i}^3 \left( J_1 |\dot{\gamma}_i| + J_2 \sum_{j \neq i} |\dot{\gamma}_j| \right) + \frac{K}{L_{d,i}^3} \quad (2.5)$$

with  $J_1$  the self hardening coefficient,  $J_2$  the latent hardening coefficient, and  $K$  a temperature dependent parameter describing dislocation recovery.

### 2.1.3 Precipitation/solid solution hardening

The following derivation presents a general form the model. The model tracks a number of precipitate phases indexed by  $i$ . The model describes these phases with the average radius  $r_i$ , the number volume density  $N_i$ , and the volume fraction  $f_i$ . Only two of these free parameters are independent but the model evolves the rate equations for all three to help with numerical stability. The growth of each precipitate phase is in turn controlled by the concentration of one or more key atomic species in the solid solution, indexed by  $j$  with mass fractions  $c_j$ . However, the chemical concentrations for non-competing precipitation reactions (see below) are directly related to the precipitate volume fractions, so the model only maintains three internal variables per precipitate phase:  $r_i$ ,  $N_i$ , and  $f_i$ .

The model represents two precipitation growth regimes. The first stage is diffusion-controlled Zener growth. Here is the concentration of the relevant species in solution are above their equilibrium concentrations for the current metal temperature and they will diffuse out of solution into the precipitates. This stage continues until one of the species contributing to the reaction reaches its thermodynamic equilibrium concentration in solution. At this

point no new precipitates can nucleate but the existing precipitates can continue to grow with larger precipitates absorbing smaller precipitates through Ostwald ripening.

In the growth regime the slowest diffusing species limits the nucleation and growth rates. The species that reaches its thermodynamic equilibrium concentration first controls growth saturation and the switch to the ripening mechanism. The particular model for 316H only tracks the concentration of Mo for the Laves phase, as there is ample iron, but tracks both C and Cr for the  $\text{Cr}_{23}\text{C}_6$  carbides. For the carbides the diffusion of Cr limits the growth and nucleation rate but the concentration of C controls the switch to Ostwald ripening.

The internal variables in the final model are coupled in two ways:

1. During the growth phase precipitation pulls the relevant species out of solution and into the precipitates. This increases precipitation hardening but decreases solid solution strengthening. Precipitate growth is therefore directly coupled to the element volume fractions in solution and all atomic species contributing to a particular precipitation reaction are likewise coupled.
2. The model sums up the contributions of all precipitate phases into a single slip resistance for precipitation hardening. Likewise, the model sums up the solid solution strengthening contributions of all tracked atomic species in the solution.

In theory, the same element in solution could contribute to multiple precipitation reactions and those reactions would compete for in pulling that element out of solution and into precipitates. This does not occur for the particular 316H model but could be added to the model relatively easily to support other materials.

The precipitate hardening strength is defined as

$$\tau_p = \frac{\alpha_p G b}{L_p} \quad (2.6)$$

with  $\alpha_p$  the interaction coefficient and

$$L_p = \sqrt{\frac{1}{\sum_i 2r_i N_i}}. \quad (2.7)$$

The solid solution strength is defined as

$$L_s = \sqrt{\frac{1}{b \sum_j \frac{c_j}{v_{m,i}}}} \quad (2.8)$$

where  $v_m$  is the molecular volume of the precipitate phase.

In the current implementation the precipitation reactions are independent from one another. As such, the subsequent exposition drops the index  $i$  representing each precipitate phase. The final model maintains and evolves one set of internal variables according to the subsequent system of coupled ordinary differential equations for each precipitate phase.

The chemical concentrations of the species contributing to each precipitation reaction are linearly related to the precipitate volume fraction:

$$c_j = \frac{c_{0,j} - f c_{p,j}}{1 - f} \quad (2.9)$$

with  $c_{0,j}$  the initial concentration in solution and  $c_p$  is the concentration in the precipitate. The dominate mechanism switches from growth to ripening based on the current volume fraction of the key species, i.e. the model is in the growth regime when

$$c_j < c_{eq,j} \quad (2.10)$$

with  $c_{eq,j}$  the equilibrium concentration for all species contribution to a given precipitation reaction. The model is in the growth regime when

$$c_j = c_{eq,j}. \quad (2.11)$$

The ODEs controlling the growth and nucleation of each precipitate then have two separate rates: one for the growth phase and a second for the ripening phase.

Finally, the Gibb's free energy driving precipitate growth is

$$G_v = -\frac{kT}{v_m} \ln \frac{c_{eff}}{c_{eff,eq}} \quad (2.12)$$

with

$$c_{eff} = \prod_j c_j. \quad (2.13)$$

The temperature sensitivity of the controlling diffusivity for each reaction follows an Arrhenius model

$$D = D_0 \exp\left(\frac{-Q_0}{RT}\right) \quad (2.14)$$

with  $D_0$  the reference diffusivity,  $Q_0$  an activation energy, and  $R$  the gas constant.

Given this setup, the evolution equation for the precipitate radius in the growth phase is

$$\dot{r}_{growth} = \frac{D}{r} \frac{c_j - c_{eq,j}}{c_{p,j} - c_{eq,j}} + \frac{\dot{N}_{growth}}{N} (r_c - r) \quad (2.15)$$

where

$$r_c = -2 \frac{\chi}{G_v} \quad (2.16)$$

with  $\chi$  the interface energy,  $\dot{N}_{growth}$  the nucleation rate in the growth regime, defined below, and *the concentration  $c_j$  is the concentration for the diffusion rate-limiting species in the reaction.* In the ripening regime the radius growth rate is

$$\dot{r}_{ripening} = \frac{M}{3r^2} \quad (2.17)$$

with

$$M = C_f \frac{8\chi V_m D c_j}{9RT} \quad (2.18)$$

where  $V_m$  is the molar volume,  $V_m = N_a v_m$  with  $N_a$  Avogadro's number, and  $C_f$  is a coarsening reduction factor linked to an increase in solubility of C in the matrix linked to the Laves phase growth.

The nucleation rate (i.e. the rate of change of the number volume density) in the growth regime is

$$\dot{N}_{growth} = N_0 Z \beta \exp\left(-\frac{G^*}{kT}\right) \quad (2.19)$$

with

$$G^* = \frac{16\pi\chi^3}{3G_v^2} \quad (2.20)$$

and

$$Z\beta = \frac{2v_m D c_j}{a_m^4} \sqrt{\frac{\chi}{kT}} \quad (2.21)$$

with  $c_j$  again for the rate limiting species and  $a_m$  the FCC lattice parameter. In the ripening regime the nucleation rate is

$$\dot{N}_{ripening} = -\frac{3N}{r} \dot{r}_{ripening}. \quad (2.22)$$

Finally, for both the growth and ripening regimes the volume fraction evolution equation is

$$\dot{f} = \frac{4}{3}\pi \left( \dot{N} r^3 + 3N r^2 \dot{r} \right) \quad (2.23)$$

which is simply the rate of the simple geometric relation

$$f = \frac{4}{3}\pi r^3 N. \quad (2.24)$$

In the ripening regime the nucleation rate is negative and the radius growth rate is positive, implying that the larger precipitates cannibalize the smaller precipitates. The overall volume fraction rate is zero, meaning the total precipitate volume remains constant.

Three types of parameters define the model, given the Tables 2.1-2.3:

1. Basic physical constants which would not vary with heat-to-heat material variation in the material structure (Table 2.1)
2. Mechanism-specific kinematic quantities like diffusivities, activation energies, etc. which may vary somewhat heat-to-heat but likely not significantly (Table 2.2).
3. The initial conditions for the internal variables, which in turn controls the initial dislocation density, precipitate distribution, and solid solution chemistry (Table 2.3). These quantities certainly vary between heats of materials and would be the primary target for calibrating a physically-based model to explain variations in the resulting creep rupture test data.

The unit system is compatible with units of megapascals for stress, Kelvin for temperature, and seconds for time.

In the current implementation the initial carbide and Laves phase radius and volume fraction are set to essentially initialize the model to have zero of either phase precipitated at the start of the simulation. This could easily be adjusted to account for the initial carbides present in 316H.

Parameter	Description	Phase/element	Value	Units
$k$	Boltzmann constant	n/a	$1.3806485 \times 10^{-23}$	J/K
$R$	Universal gas constant	n/a	8.31462	J/(K · mol)
$N_a$	Avogadro's number	n/a	$6.02 \times 10^{23}$	1/mol
$a_m$	FCC lattice parameter	n/a	$3.6 \times 10^{-10}$	m
$b$	FCC Burgers vector	n/a	$2.5 \times 10^{-10}$	m
$\alpha_0$	Average interaction coefficient	n/a	0.5	n/a
$c$	Energy barrier shape constant	n/a	3/4	n/a
$d$	Energy barrier shape constant	n/a	4/3	n/a
$\dot{\gamma}_0$	Reference slip rate	n/a	1	1/s
$G_0$	Shear modulus at absolute zero	n/a	$77000 \times 10^6$	Pa
$G$	Shear modulus at temperature	n/a	$773.15\text{ K}$ 823.15 K 873.15 K 923.15 K	MPa
$V_m$	Molar volume	Carbide Laves	$6 \times 10^{-6}$ $2 \times 10^{-6}$	m <sup>3</sup> /mol
$C_p$	Concentration in precipitate	Cr in carbide C in carbide Mo in Laves	69.85 5.13 50	wt%
$C_{eq}$	Equilibrium concentration in solution	Cr C Mo	15.64 $7.25 \times 10^{-6}$ 0.25	wt%

Table 2.1: Physical constants that do not vary heat-to-heat.



Parameter	Description	Phase/element	Value	Units
$J_1$	Self hardening constant	n/a	$5 \times 10^{14}$	$1/\text{m}^2$
$J_2$	Latent hardening constant	n/a	$2.5 \times 10^{15}$	$1/\text{m}^2$
$K$	Dislocation recovery constant	n/a	$1.5 \times 10^{-32}$	$\text{m}^4/\text{s}$
$\alpha_d$	Dislocation interaction coefficient	n/a	0.35	n/a
$\alpha_p$	Precipitate interaction coefficient	n/a	0.84	n/a
$\alpha_s$	Solid solution coefficient	n/a	0.00457	n/a
$N_0$	Nucleation site density	Carbide Laves	$1 \times 10^{13}$ $5 \times 10^{14}$	$1/\text{m}^3$
$\chi$	Interface energy	Carbide Laves	0.3 0.25	$\text{J}/\text{m}^3$
$D_0$	Reference diffusivity	Cr/Carbide Mo/Laves	$1.5 \times 10^{-4}$ $7.4 \times 10^{-4}$	$\text{m}^2/\text{s}$
$Q_0$	Diffusion activation energy	Cr/Carbide Mo/Laves	240000 283000	J
$C_f$	Carbide coarsening factor		773.15 K 823.15 K 873.15 K 923.15 K	
		Carbide Laves	1 1 0.3 0.03 1 1 1 1	n/a

Table 2.2: Physical properties that do not vary significantly heat-to-heat.

Parameter	Description	Phase/element	Value	Units
$L_d$	Dislocation spacing	n/a	$1.0 \times 10^{-7}$	m
$r$	Precipitate radius	Carbide	$1 \times 10^{-9}$	m
		Laves	$1 \times 10^{-9}$	
$f$	Precipitate phase fraction	Carbide	$4.19 \times 10^{-16}$	n/a
		Laves	$4.19 \times 10^{-16}$	
$c$	Initial chemical concentration	Cr	16.25	wt%
		C	0.0375	
		Mo	2.33	

Table 2.3: Physical parameters that will vary significantly heat-to-heat. Note that the precipitate number density would also need to be initiated, but could be done using Eq. 2.24 based on the values of radius and phase fraction given here.

## 2.2 Implementation in NEML

The numerical implementation in NEML differs somewhat from the basic presentation of the model described above, drawn from the original sources cited in the previous section. The goal of these changes was to improve the performance of the model as implemented in CPFE simulations. Specifically, NEML integrates the system of ODEs defining the model state with a fully implicit scheme. The original authors decoupled the integration of the crystal plasticity kinematics and some of the internal variables. The fully implicit integration in NEML improves the stability of the model and simplifies the calculation of the algorithmic tangent, needed to achieve quadratic convergence in CPFE simulations.

The first change is rather than switch between the growth and ripening regimes abruptly when the critical element reaches the equilibrium concentration in solution, instead the model uses a weighted sum of both the growth and ripening rates, i.e.

$$\dot{r} = f(c_j)\dot{r}_{growth} + (1 - f(c_j))\dot{r}_{ripening} \quad (2.25)$$

and

$$\dot{N} = f(c_j)\dot{N}_{growth} + (1 - f(c_j))\dot{N}_{ripening} \quad (2.26)$$

where

$$f(c_j) = \begin{cases} \frac{c_j - c_{0,j}}{c_{eq,j} - c_{0,j}} & c_j \leq c_{eq,j} \\ 1 & c_j > c_{eq,j} \end{cases} \quad (2.27)$$

This mixes in the two mechanisms in proportion to how close the critical species concentration is to the equilibrium concentration. As the critical element approaches the equilibrium concentration in the solution the ripening mechanism becomes more dominant over the growth mechanism, until when the critical species reaches equilibrium the evolution rate is equal to the ripening mechanism rate. This modification approximates the hard switch between mechanisms in the original model but makes the first derivative of the rate equation continuous, required to achieve good convergence when integrated with a fully implicit method. This approximation also improves the numerics of the model as a whole when used in CPFE simulations.

Variable	Scale factor	Scaled units
$L_d$	$10^{-9}$	nm
$f$	0.1	No simple unit
$r$	$10^{-9}$	nm
$N$	$10^{12}$	Parts per trillion

Table 2.4: Scaling factors for the internal variables.

The second change is that the dislocation spacing and precipitate radius, number density, and volume fraction equation values are all scaled by dividing by a scale factor to improve the numerics of the model by making the values of the internal variables approximately equal. Essentially all this does is change the units used in tracking the internal variables. However, the NEML implementation does this internally, so that the user can work with the original unit system presented in the reference material and defined in Tables 2.1-2.3. Table 2.4 lists the specific scale factors used in the simulations below.

## 2.3 Verification

To verify the model we ran thermal aging simulations for 316H using the model implemented in NEML. These simulations are for a single material point representing a single crystal (orientation does not matter). The simulation holds the stress fixed to zero, the temperature fixed to one of 550°, 600°, or 650° C, and integrates the system of ODEs representing the evolution of the internal state through time, out to  $10^7$  hours. The chemical concentrations of C, Cr, and Mo and the radius, number density, and volume fraction of the carbide and Laves phases evolve with time, along with the total precipitate obstacle spacing. We can compare these values to the information plotted in [21], which presents the results of similar thermal aging simulations.

Figures 2.1-2.3 show the results for all three temperatures. The chemical concentration in the solid solution agree nearly exactly with the results presented in [21], with the only difference occurring near the transition to the equilibrium concentration. The modification described above, i.e. mixing the growth and ripening rates, explains this small difference. There are somewhat larger differences in the carbide and Laves phase volume fraction and area density results. These are partly attributable to the modifications to the numerical implementation of the model but also likely due to differences in how we integrate the rate equations (implicit here versus explicit in the original paper). All these small differences contribute to produce a somewhat different precipitate spacing versus time relation, which is what feeds into the slip system hardening. However, our implementation produces the same basic curve, just scaled to reflect the implementation differences. Overall, these differences do not end up producing a significantly different physical response, i.e. the thermally-aged flow curve simulated by embedding the model in a homogenized or full field crystal plasticity simulation.

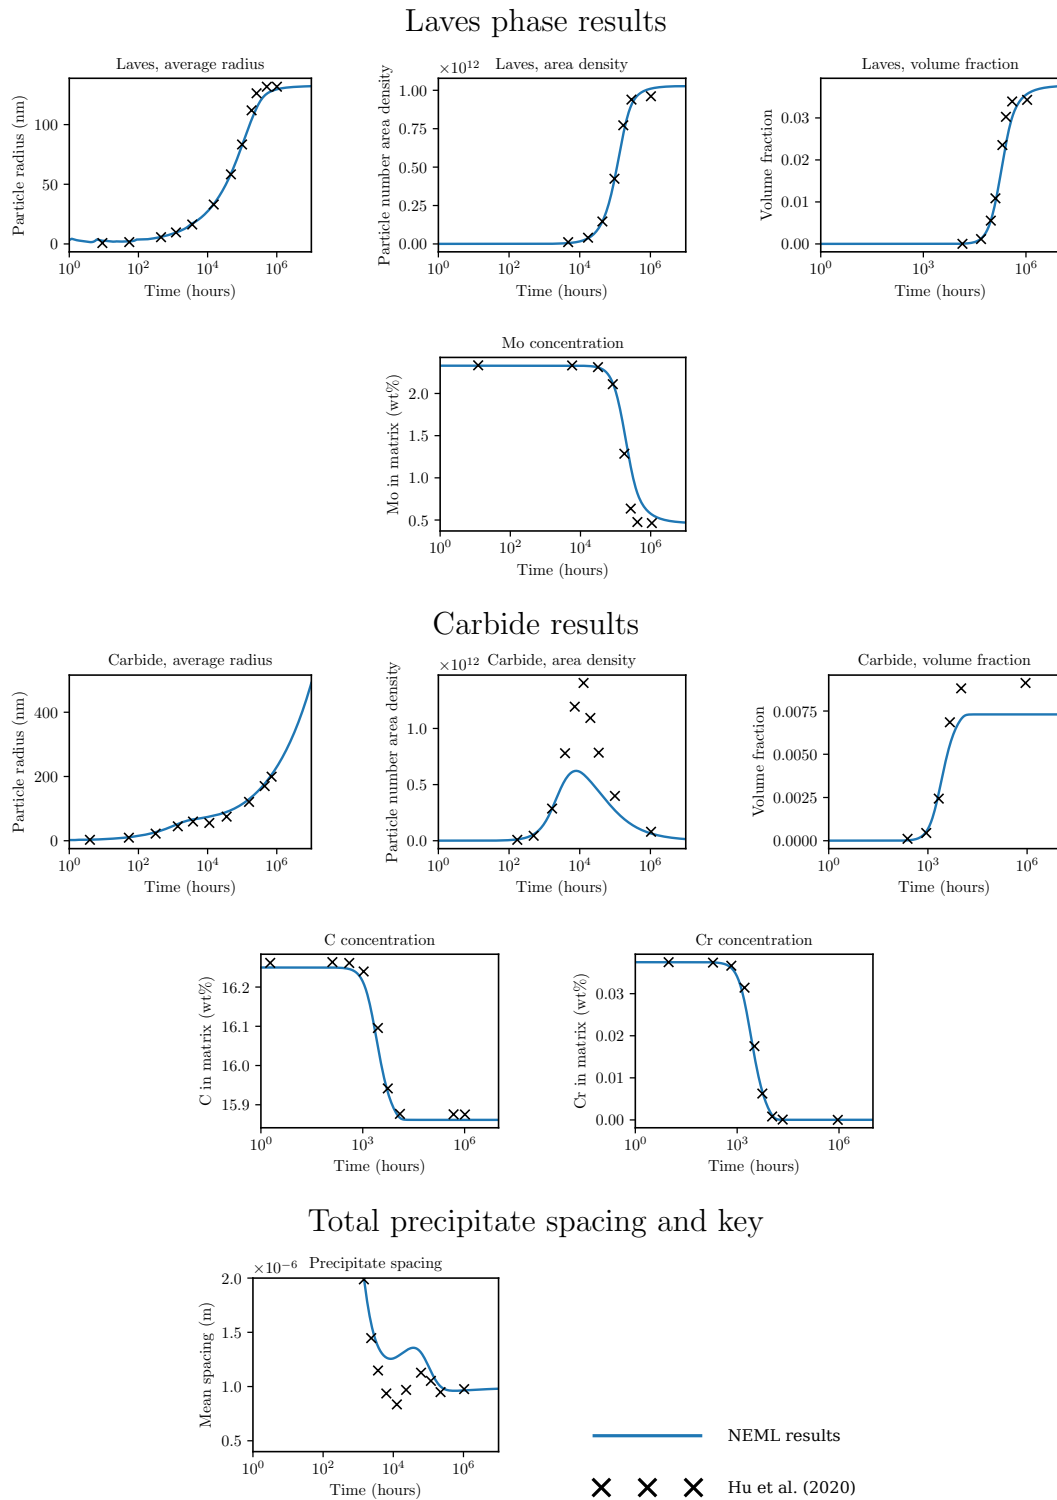


Figure 2.1: Comparison between NEML implementation and [21] results for thermal aging at  $T = 550^\circ \text{C}$ .

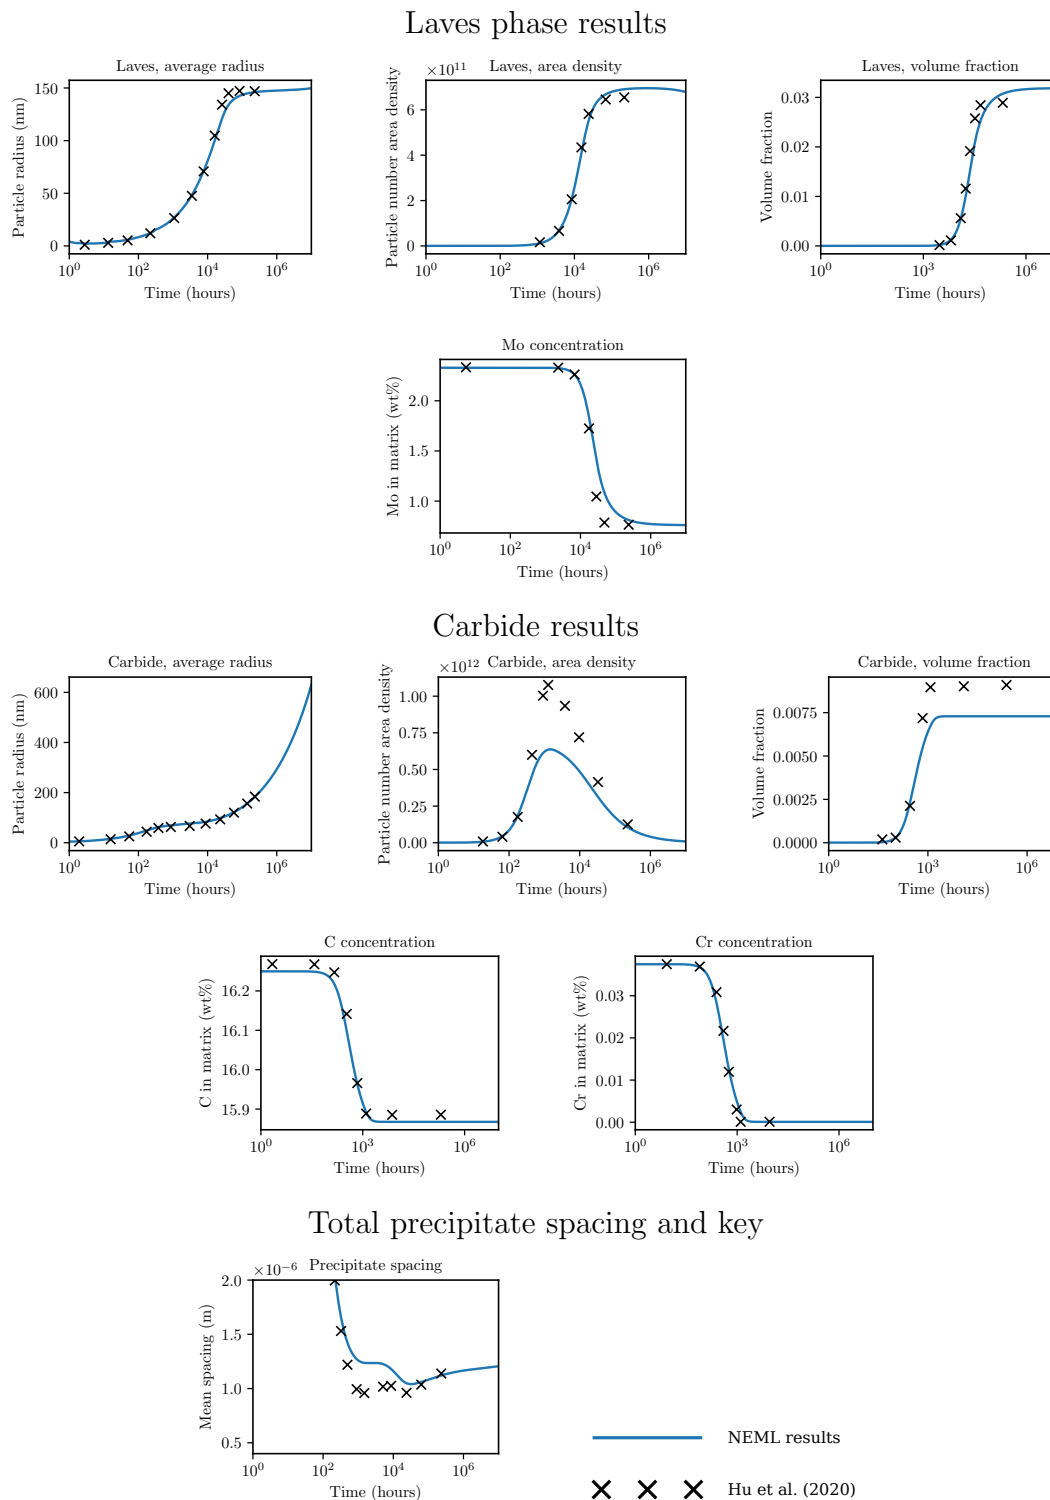


Figure 2.2: Comparison between NEML implementation and [21] results for thermal aging at  $T = 600^\circ \text{C}$ .

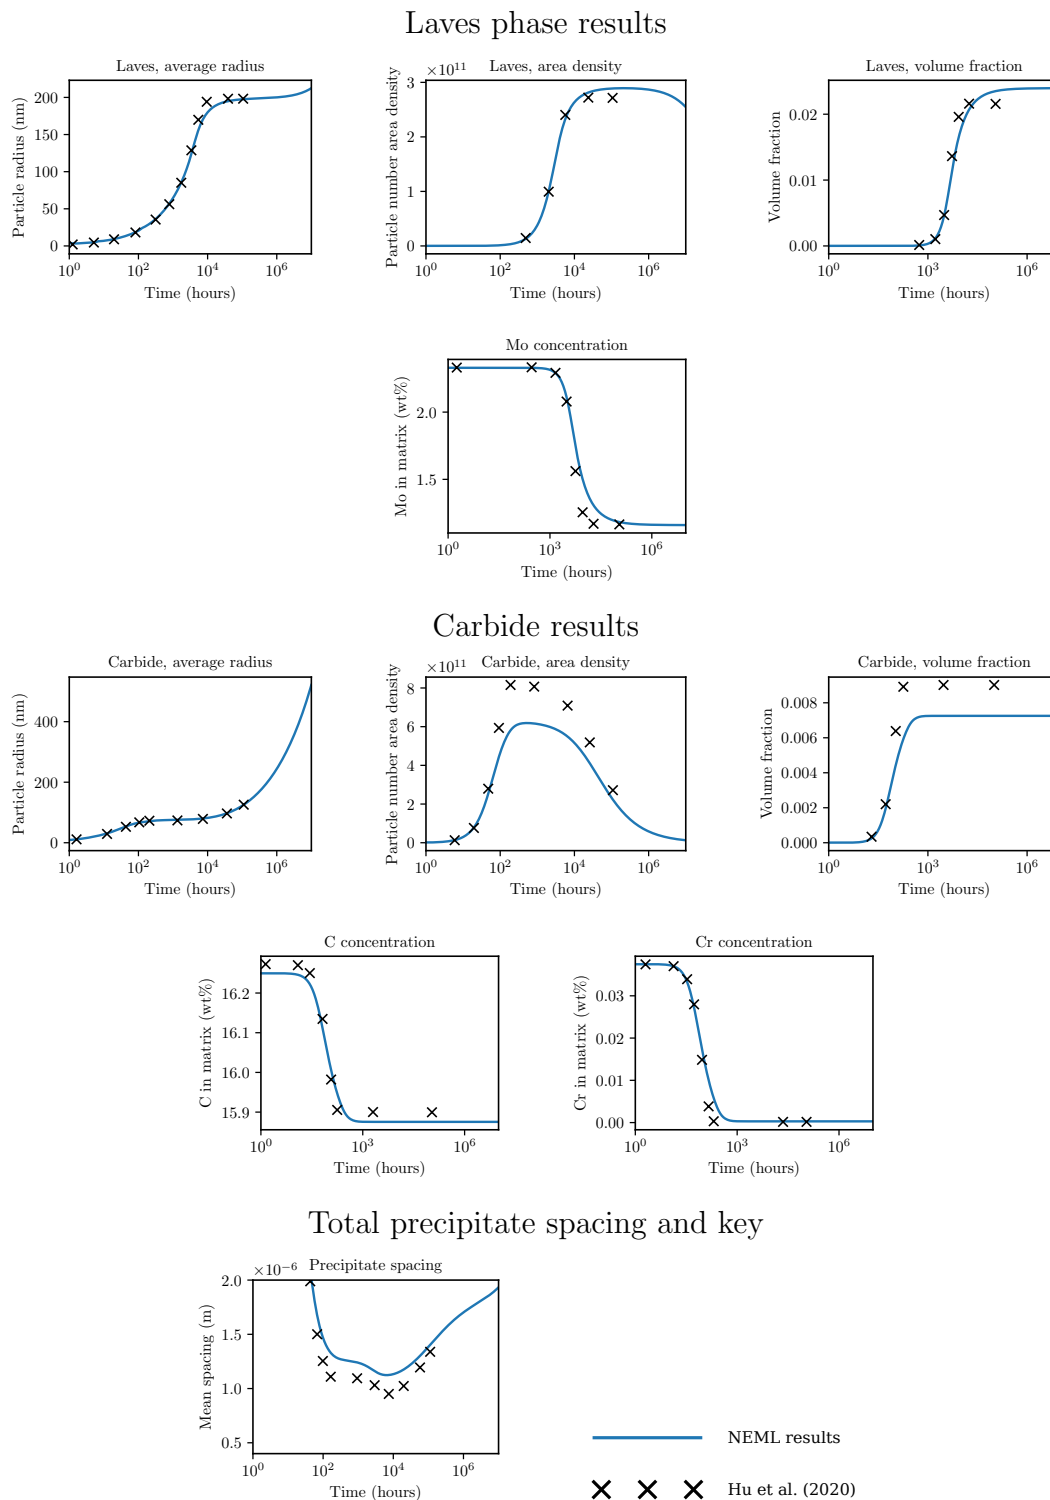


Figure 2.3: Comparison between NEML implementation and [21] results for thermal aging at  $T = 650^{\circ} \text{C}$ .




Figure 2.4: Unaged tension response of the Hu and Cocks model.




Figure 2.5: Aged tension response of the Hu and Cocks model. The solid lines are the unaged response, the dashed lines are the response after aging 100,000 hours at the test temperature.

## 2.4 Macroscale examples

This section illustrates the macroscale response of the model using polycrystal simulations. These simulations all homogenize the model response to average macroscale behavior with Taylor-homogenized simulations of 100 randomly selected crystal orientations.

Figure 2.4 shows the unaged tension model response by simulating a standard tensile test with a strain rate of  $8.33 \times 10^{-5}$  mm/mm/s at temperatures of 500°, 550°, 600°, and 650° C. The model response is reasonable for 316H at these temperatures, including both the initial yield stress and work hardening.

Figure 2.5 then repeats these tension test simulations but now after first aging the material for 100,000 hours at the test temperature. The figure compares the unaged and aged tensile response. Thermal aging does not significantly affect the flow stress of 316H for these conditions, which is as expected based on the experimentally-observed response of the material.




Figure 2.6: Simulated creep tests using the Hu and Cocks model.

Finally, Figure 2.6 plots the creep response of the model at 500°, 550°, 600°, and 650° C for a stress level of 40 MPa. Again, this figure illustrates that the model performs reasonably compared to the real material response at these conditions.

Overall, these results demonstrate the model is ready for use in CPFEE calculations. However, as described in Chapter 1, we started the CPFEE throughput simulations in parallel with the development of this new model, using a pre-existing model form already implemented in the CPFEE framework.



### 3 CPFEM model and throughput simulation database

This chapter describes a campaign of throughput crystal plasticity (CP) simulations, which in turn feed into the physics-based model for extrapolating creep rupture described Chapter 4. Crystal plasticity simulations were performed on a 100-grain representative volume element (RVE) to simulate creep. The simulations were performed to generate a database of creep rupture times for 316H steel for a temperature of 600° C. The input parameters for the CP model were varied about a baseline value; the baseline parameters values were established by comparing the model creep response to experimental creep curves. For these simulations, standard Voce-type saturation laws are used to model the hardening response of grain interiors and the grain boundary cavitation model developed in an earlier study [25] is used to model the grain boundary deformation. The throughput simulation grid varies six parameters, selected based on past sensitivity studies, specially three grain boundary parameters ( $a_0$ ,  $b_0$  and  $D_{GB}$ ) for the grain boundaries and three Voce hardening parameters ( $\tau_{sat}$ ,  $\tau_0$  and  $b$ ) for the grain bulk.

#### 3.1 Simulation setup

A 100-grain RVE microstructure was used as the input microstructure for all the simulations. Six crystal plasticity parameters are varied about their baselines values, as shown in Table 3.1, to sample the variation in the creep rupture life caused by heat-to-heat microstructural variations in the material. Figure 3.1 shows the microstructure and the loading conditions used in these simulations. The baseline parameters are used as input to the CP model, and after confirmation that the resulting creep curves are comparable to the corresponding experimental creep curves, the parameters are varied and subsequent throughput simulations are performed. The creep curves are generated for 5 different stress levels: 20 MPa, 50 MPa, 100 MPa, 150 MP and 200 MPa, and at 600 °C. Figure 3.2 shows plots of strain and strain rates for the 100 MPa input stress case. The simulated creep curves are generally centered on the experimental curves but the range of results in the throughput simulation database span a few orders of magnitude in either direction of the experiments. This variation is much wider than what we expect for actual heat-to-heat variations in 316H. Running the simulations to cover a wider range of parameter space than what we will actually use in predicting rupture ensures that a surrogate model, trained against this data, can capture *at least* the full range of responses needed to explain the real heat-to-heat variation.

Table 3.1: The crystal plasticity model parameters varied in the throughput simulations to generate a database of creep rupture data.

Symbol	Description	Value and	Units
$a_0$	initial cavity half radius	$5 \times 10^{-5} \pm 50\%$	$mm^2$
$b_0$	initial cavity half spacing	$0.06 \pm 50\%$	$mm^2$
$D_{GB}$	GB diffusion coefficient	$1 \times 10^{-15} \pm 50\%$	$mm^3/MPa.h$
$\tau_{sat}$	Isotropic saturation	$12 \pm 50\%$	MPa
$\tau_0$	Isotropic initial value	$40 \pm 50\%$	MPa
$b$	Voce hardening pre-factor	$66.66 \pm 50\%$	unitless





Figure 3.1: (a) The 100-grain microstructure used for all the simulations to generate the database of creep rupture times. (b) The loading conditions imposed on the microstructure.





Figure 3.2: Distribution of creep (a) strain and (b) strain rate curves from the different throughput simulations, compared with the experimental creep curve for the same loading conditions. These curves are all for a stress of 100 MPa.

## 3.2 Analysis

The simulations produce full creep curves as well as the detailed microstructural state of the material as a function of time. From these results we calculate a surrogate for the creep rupture time to produce a database linking the stress and the six key microstructural parameters to the resulting time to rupture. To account for the large variations in the behavior across the parameter space and different stress levels, the creep rupture time is calculated as the time required for creep cavitation to increase the simulation cell volume by a factor of 0.01. First, the volume increase due to elasticity is subtracted from the total change in volume. Since the creep load begins at  $t = 0.1hrs$ , all volume change before that is attributed to elasticity and ignored. Then, the time corresponding to a volume increase of 0.01 is assigned as the creep rupture time. Two examples are shown in Figure 3.3, which illustrate the calculation of the creep rupture times as an increase in the excess volume of the microstructure. These two cases represent brittle and ductile creep failure. A relatively low threshold of 0.01 is used for the increase in excess volume to account for the large variation in the creep behavior across different stress values and input microstructural parameters. In some cases, we do not reach the chosen threshold within the simulated time frame, and so we do not include these simulations in the subsequent analyses.

The rationale behind this surrogate measure of rupture is to avoid attempting to define a definite rupture time by pulling the simulation volume in half. The simulation results become very sensitive to the details of the numerical implementation of the grain boundary and grain bulk constitutive models near the end of life when tertiary creep causes large creep strain rates and accumulated strains. This sensitivity makes it difficult to assign a consistent time-to-failure to a wide range of simulations. As the failure mechanism in these simulations is grain boundary cavitation, an increase in volume due to cavity growth is a reasonable surrogate for damage. More importantly, this surrogate measure respects the differences between microstructural parameters causing a creep brittle response versus parameters causing a creep ductile response. In effect, the cavitation measure is a reasonable way to indicate failure that does not limit the available creep ductility. For more details on this criterion see [26].




Figure 3.3: Calculation of the increase in the excess volume (after elasticity) for (a) brittle creep failure and (b) ductile creep failure cases.

## 4 Bayesian inference model and long-term predictions of creep rupture

This chapter describes the process of taking the results of the CPFE throughput simulations described in Chapter 3, using actual, short-term creep rupture test data on 316H to infer a statistical model for the creep rupture life of the material, and then using this statistical model to predict the long-term creep rupture strength of the material. As described in the introduction, because long-term rupture test data is available for 316H steel we can then evaluate the accuracy of this approach against actual test data. This chapter additionally compares the accuracy of the extrapolated model against a Larson-Miller model, assessing our new physics-based approach to the current, empirical methods used to qualify materials.

For the example here we use creep rupture data for 316H at 600° C. We selected this temperature because it represents a realistic operating temperature for future advanced reactor components and because the existing rupture database has a large number of tests at this temperature, including tests with rupture times greater than 200,000 hours. The results here train the final model using the portion of the rupture database with failure times less than 15,000 hours (less than 2 years) to evaluate the accuracy of the approach in extrapolating based on a short-term, accelerated test program.

### 4.1 Surrogate modeling

The overall goal of the modeling work (see Fig. 1.3) is to use Bayesian inference to find the distributions of the physical, microstructural parameters input to the CPFE model that explain the resulting distribution of actual creep rupture data. However, the CPFE framework is too numerically expensive to use directly in the inference process. Instead, as described in the previous chapter, we assembled a large throughput simulation data of CPFE results. This throughput database covers the stress range of interest for practical component designs as well as sampling a wide range of the parameter space for the six key microstructural parameters. The database maps these seven parameters to the resulting simulated critical cavitation time, a surrogate for time to rupture under creep conditions.

The goal here is to fit a surrogate model to this data to replace the CPFE model in the inference process. This model should have the same input/output as the CPFE results, i.e. it maps the stress and the six microstructural parameters to the critical cavitation time, but needs to be much less computationally expensive than a direct CPFE calculation.

The total database includes the results of 1651 calculations. Before fitting a surrogate model we culled the CPFE database to eliminate failed or unreliable simulations results. Some simulations did not reach the critical void fraction over physical time allocated for the simulation. We eliminated these from the database, leaving 1607 simulations. We then culled all the simulation data with predicted critical cavitation times greater than 900,000 hours. There are two reasons to remove these data:

1. They are unnecessary as there is no real need to predict rupture times of this length — it is very unlikely that any component in a future reactor will operate for times greater than 100 years, which is approximately 900,000 hours.
2. Some of the simulations for very low stress conditions, i.e. very long rupture times, produced anomalous creep curves where creep deformation essentially stopped for long




Figure 4.1: Summary histogram of the dataset after filtering.

periods of time. Likely this is a numerical issue related to the order of magnitude of the resulting creep rates: very low stress values produce very slow creep rates, which tend to fall below the integration tolerances used by default in the CPFE calculations. This could, at least potentially, be rectified either by adjusting the integration tolerances or by switching the units of the calculations from seconds to hours. However, given that these long times are not very relevant to engineering design we feel justified in simply excluding these data.

Figure 4.1 summarizes the dataset after filtering. The plot shows a histogram of the number of points at each temperature and the resulting time to failure. Culling the dataset as described above most affects the 20 MPa stress simulations, as these generally have the longest rupture time. The dataset still contains 1340 simulations after culling the long rupture time data.

We used a deep neural network trained to this data as a surrogate model to replace direct CPFE simulation. The network has an input depth of 7, i.e. the stress plus the 6 microstructural parameters. It has two intermediate layers of size 10 and 5 and an output size of 1 (i.e. the critical cavitation time). The network applies an ReLU activation function at each layer. We implement the model and the calibration process in the pytorch [27] framework.

In training the model we first randomly split the filtered dataset into training and test sets. Specifically, the process randomly selects 10% of the data to be in the test set and the remaining 90% for the training set. We use pytorch to train the neural network parameters against the data using the Adam gradient-descent [28] optimizer with a learning rate of




Figure 4.2: Typical training history for fitting the surrogate model to the CPFE data.

$5 \times 10^{-4}$  over 15,000 iterations.

Before training we first take the base-10 logarithm of the stress (input) and the rupture time (output data). Moreover, we shift and scale the input data, both the stress and the microstructural parameters, so that they are centered about zero and have a scale of 1. The loss function is the mean squared error between the surrogate model predictions and the actual CPFE time to rupture. Arguably, a relative error measure might perform somewhat better as the mean squared error tends to weigh the longer rupture data more than simulations with short simulated rupture times. However, taking the log of the times before training minimizes this effect and the final model performs adequately well for all rupture times.

Because the initial values of the neural network weights and biases are random the final model is somewhat stochastic. A section below describes the variation in the final results for multiple repetitions of the entire process described in this chapter. Figure 4.2 plots the convergence history of the training process, which is representative of any repetition of the surrogate model calibration process.

Figure 4.3 illustrates the accurate of the model by plotting the surrogate model versus CPFE rupture time *for the reserved test dataset*. The results, as well as the actual error calculated over the reserved data, show that the model is generally accurate. Using the reserved test demonstrates that the training process is robust against overfitting. The surrogate is generally quite accurate and, moreover, the error between the surrogate model and the test data concentrates at very long rupture times, which are not relevant for realistic component designs.




Figure 4.3: Plot comparing the surrogate-predicted and actual CPFE rupture times, using the reserved test data.

## 4.2 Bayesian inference

The surrogate model provides a rapid means to map the stress and microstructural parameters to the CPFE-predicted rupture time:

$$t_R = f(\sigma, \mathbf{m}; \mathbf{p}) \quad (4.1)$$

where  $\mathbf{m}$  is the vector of the six microstructural parameters and  $\mathbf{p}$  are the neural network parameters, trained against the CPFE data.

The goal of the inference is to find the distribution of the microstructural parameters that predicts the distribution in creep rupture life over the short-term test database. Mathematically,

$$T_R = f(\sigma, \mathbf{M}) \quad (4.2)$$

where  $T_R$  and  $\mathbf{M}$  are now random variables but  $\sigma$  remains deterministic. More importantly, the internal neural network parameters are not included in the inference process. They remained fixed, deterministic parameters set to the values determined by the surrogate modeling fit against the CPFE data. This approach in effect *replaces* the CPFE simulations with the surrogate model for determining the microstructural parameter distributions using the actual rupture data. By keeping the internal surrogate parameters the same, we are in effect fixing the functional form of the model to match the physics-based CPFE simulations. An alternative approach using a Bayesian neural network would find probability distributions for these internal parameters based on the actual data. This process would lose the information encoded in the CPFE results. Because we retain the network structure unaltered and only infer the microstructural parameter distributions our final Bayesian model is physically based – constrained to behave as if we are inferring microstructural parameter distributions directly through the CPFE calculations.



The statistical model used here also accounts for random experimental and measurement variations by including a white noise term:

$$T_R = f(\sigma, \mathbf{M}) + \epsilon \quad (4.3)$$

where  $\epsilon$  is Gaussian noise with zero mean. This noise term accounts for variations that cannot be modeled as heat-to-heat microstructural variations, e.g. variation in the rupture life for samples from the same heat. These types of variations are common in creep tests.

We use a large collection of 316H rupture data to calibrate the statistical model [20]. For the example here, we set a  $t_{limit} = 15,000$  hours, of somewhat less than 2 years. The inference process includes all the data with rupture times less than this limit. After inferring the model we can then compare the extrapolated predictions of creep rupture time against the data with rupture times greater than  $t_{limit}$ . At a temperature of  $T = 600^\circ \text{C}$  the rupture database contains 367 tests. For a limit of 15,000 hours there are 286 tests with rupture times less than the limit and 81 tests with rupture times greater than the limit.

We use the variational Bayes inference method [19] as implemented in the Pyro python library [29]. The complete details of the method are contained in these references. The objective of the approach is the same as for any Bayesian inference approach: to estimate the posterior distribution of the model parameters that best explains the data. As with other inference techniques, like Markov Chain Monte Carlo (MCMC), a set of prior distributions provided by the user informs the inference method by providing the best available information on each parameter distribution. Unlike MCMC approaches, variational Bayesian techniques approximate the posterior distribution with some parameterized, analytical probability distribution. In pyro this distribution is called the guide. The inference approach then optimizes the variables parametrizing this guide distribution to best explain the experimental data. Specifically, pyro maximizes the evidence lower bound (ELBO).

We structure the statistical model as a hierarchical, pooled description of the heat-to-heat variation of interest here (Fig. 4.4a). First, the model samples two random variables for each of the six microstructural parameters. These “top-level” random variables represent the location and scale of the parameter distributions describing variation over randomly selected heats of material. Then, for each experiment, the model samples the actual microstructural parameters from the distributions described by the top-level variables. This two-level model exactly describes heat-to-heat variation under the assumption that each experimental sample is drawn from a random heat. This process could be improved if batch information was available for each test by instead sampling the top-level parameter distributions once for each heat of material, instead of once for each sample, and simulating the rupture time for all samples in the same heat using the same set of heat-specific microstructural parameters (Fig. 4.4b).

For variational Bayes techniques the user must provide then both the prior distributions and the form of the guide distribution used to approximate the posterior distribution. We use a normal prior for the top-level random variables describing the parameter locations and a half normal prior for the top-level random variables describing the parameter scales. As the parameters of the physics-based model all have a physical meaning, these prior distributions would ideally be informed by microstructural characterization. However, for now we set the mean of these distributions to the average of the parameter grid used to run the throughput




Figure 4.4: Figure illustrating the concept of a hierarchical, pooled model to capture heat-to-heat variation in the model microstructural parameters.

calculations. This represents some slight information, as we ensured that these average parameter values produced reasonable, if not calibrated, creep curves for 316H. We set the standard deviation of the top-level location priors as well as the standard deviation of the half normal priors describing the distribution variance and the white noise to 0.2. This is mostly uninformative, set just to be larger than the actual variance of the true posterior distributions. This reflects our observation that variational Bayesian methods, at least as implemented in pyro, are much better at narrowing a wide prior than they are at broadening a narrow prior.

We use a delta distribution for the guide distributions corresponding to the top-level random variables describing the heat-to-heat microstructural parameter distributions. This means that the method will not find the distribution of these top-level parameters but instead a pointwise maximum a posteriori (MAP) estimate of the parameter distribution locations and scales. In other words, setting up the inference model this way provides a point estimate of the location and scale of the microstructural parameter distributions. This is precisely what we want. The MAP estimate is still regularized by the prior information and so this approach combines the best of a Bayesian framework with the goal of finding point estimates of the metaparameters describing the heat-to-heat parameter variations.

Figure 4.5 shows a typical training history using the Adam optimizer [28] with a learning rate of  $5 \times 10^{-4}$ . A good amount of the complexity of the underlying the pyro implementation of variational Bayesian inference is setting up the problem to provide an unbiased, low variance estimate of the gradient of the loss with respect to the guide parameters [30]. The loss function is now by definition stochastic as evaluating the ELBO involves sampling the heat-to-heat distributions for each experiment. We draw four Monte Carlo samples each iteration to minimize the effects of random sampling, and despite this variation on average




Figure 4.5: Typical loss history for the inference optimization process.

the loss decreases each iteration to a minimum.

Figure 4.6 plots the normalized heat-to-heat parameter posterior distributions for one repetition of this inference process. The parameters are normalized so that the prior mean is zero. The inference process considerably narrows the variance of the parameter distributions compared to the priors and shifts the means somewhat to match the experimental data. Without better prior information there is no real guarantee that the inferred microstructural parameter distributions match the actual material. All we can say is that the resulting microstructural parameter distributions are physically reasonable.

Figure 4.7 summarizes the key results of the study – predicting the distribution of the long-term creep rupture strength of 316H using only short term data. This plot compares the data, the physics-based Bayesian model described here, and a linear Larson-Miller correlation, fit using the same data available to the inference model. Specifically, this plot shows:

1. The experimental data used to calibrate the models.
2. The long-term experimental data not included in the model calibration.
3. The mean predicted life for the Bayesian model and the linear Larson-Miller model.
4. A 95% prediction interval for both models.

Already this plot demonstrates that the physics-based model can capture long-term behavior the empirical model cannot capture. The inference model correctly predicts a shift in the slope of the rupture correlation for both the very short-term/high-stress data (included in the training data) and, critically, for the long-term extrapolation predictions. These mechanism shifts reflect the underlying physical behavior of the material, likely the effect of substantial dislocation hardening for the high stress tests (where the creep test stress




Figure 4.6: Example normalized posterior distributions for the microstructural parameters inferred from the short-term rupture data.




Figure 4.7: Comparison between experimental data, the physics-based model, and a linear Larson-Miller model, plotting predicted rupture time as a function of the applied stress.

is greater than the material yield stress) and a shift towards diffusional creep mechanisms for the long-term predictions. The physics-based model extrapolation predictions are more accurate than the empirical Larson-Miller model, both in terms of the mean trend lines and in how many of the actual experimental rupture times were captured by the 95% prediction bound.

### 4.3 Assessing the results

One drawback of the physics-based Bayesian model is that the results of the complete process going from CPFE data to inferred statistical model are somewhat non-deterministic. This random variation arises from several sources:

1. The initial split of the CPFE data into the training and test sets is random.
2. The optimization process for fitting the surrogate model starts from random initial network scale and bias parameters.
3. The variational Bayesian inference process is itself stochastic. The initial guide parameter values are sampled randomly from the priors and the loss function and gradient themselves are stochastic.

As such, to evaluate the model accurately we repeated the entire process, from CPFE data all the way to extrapolated rupture time predictions, 10 times. The objective of this work is to predict long-term creep failure based on short-term data and so we focus on two error metrics assessing the models against the long-term creep data excluded from the training database:

1. The mean absolute relative error between the average model prediction and the actual rupture data.
2. The number of points in the prediction region that actually fall inside the model 95% prediction bound.

The first metric assesses the general accuracy of the approach. The second metric would be crucial for developing design material properties, as these are typically based on a minimum, lower-bound (often indeed a 95% interval) on experimental measures. Table 4.1 summarizes the results for the inference model and both a linear and quadratic Larson-Miller correlation. The Larson-Miller results are deterministic – there is no random element in fitting these models. The table provides both the mean metrics over 10 repetitions of the complete process for developing the physics-based model as well as the maximum and minimum values from all 10 repeats. This interval provides a sense for how repeatable the results are as well as how safe it would be to make predictions based on a single instantiation of the model.

In all cases and for both the mean relative errors and % correctly predicted within the 95% prediction interval the physics-based Bayesian model outperforms the empirical linear and quadratic Larson-Miller extrapolations. While the error metrics do vary somewhat for multiple repetitions of the calibration process, the variation is not significant. A perfect model for developing design values of creep rupture strength would correctly predict the

		Mean relative error		Within 95% interval	
Linear Larson-Miller		10.2%		83.9%	
Quadratic Larson-Miller		11.9%		77.8%	
Physics-based	Min		6.4%		87.6%
	Mean	8.2%		90.0%	
	Max		9.6%		91.4%

Table 4.1: Comparison of the accuracy of the physically-based Bayesian model as well as linear and quadratic Larson-Miller correlations in extrapolating long-term creep rupture data based on short-term data. For the Bayesian model, the table shows results from 10 repetitions of the complete process.




Figure 4.8: Relative error versus long term,  $t > 100,000$  test data for the physics-based model with different values of  $t_{limit}$  (i.e. the longest duration creep test kept in the inference database).

true 95% interval on the long-term data. However, even conventional Larson-Miller models calibrated against long-term experiments meet this criteria when extrapolating. Determining whether the approximate 90% accuracy is sufficient will at least partly be a matter of engineering judgement, although future analysis may compare the accuracy of Larson-Miller models calibrated over the complete, long-term database to extrapolated predictions with the physics-based model to provide some insight into what level of accuracy should be expected.

Finally, we can use this model to examine the acceptable extrapolation factor: what length of creep rupture data is needed to correctly predict long-term rupture strength on the order of the design lives of future components. Figure 4.8 plots the accuracy of the physics-based model in predicting the rupture life for the very long-term data with actual, experimental rupture times greater than 100,000 hours (about 11 years, 6 data points) as a function of the time limit  $t_{limit}$ . That is, this study examines the accuracy of the model for different time extrapolation factors. This plot only shows the mean absolute relative error as there are insufficient data points to meaningfully compare to a prediction interval. For the physics-based model, the plot shows the error metric averaged over five repetitions of the model calibration process to minimize the differences between runs.

The figure demonstrates that the long-term property predictions remain accurate even if the test database only spans 5,000 to 10,000 hours. This suggests the new approach developed here can accurately predict long-term creep rupture for realistic component design lives of greater than 100,000 hours based on creep test data with maximum durations of less than one year.





## 5 Conclusions and future work

### 5.1 Summary

This report describes the development and testing of a new extrapolation method for time-dependent material properties aimed at accelerating the qualification of new materials by reducing the length of the test program required to accurately extrapolate key time-dependent material properties out to expected 30 or 60 year component design lives. While many factors contribute to the amount of time required to qualify a new material for nuclear use, the requirement for a dedicated long-term test program, particularly for high temperature reactors, is one of the most difficult to overcome.

The method described here uses a physics-based model to predict the creep rupture life of the material. The report examines two physically-based models for creep deformation and failure in 316H stainless steel: the Hu and Cocks model and a simpler model developed at Argonne National Laboratory. We generate a large throughput simulation database with the simpler model to calculate the expected rupture time for a wide range of stresses and key microstructural parameters.

The key concept described in this report is using Bayesian inference to then calibrate this physics-based model against short-term creep test data. This inference process requires less model development time, experimental characterization, and computational expense than a full multiscale approach while at the same time retaining a direct connection between the model predictions and the underlying microstructural mechanisms. In the context of Bayesian inference the prior distributions provide a means for incorporating any available microstructural characterization data into the model. The advantage of the approach is that where such information is not available the user can provide an uninformed prior and the inference process will still attempt to find the distribution of microstructural parameters that best explains the data.

Our hypothesis is that using a physics-based model, even when calibrated against short-term test data, will better predict the long-term creep rupture strength of the material. We tested this hypothesis using a large creep rupture database for 316H stainless steel. Because this already-qualified material has been extensively studied there are actual creep rupture test data for very long times — greater than 200,000 hours for the conditions considered in this report. This means we could test the models by providing a short-term subset of this data for calibration and then testing the resulting predictions against the remaining long-term data.

This comparison proves our hypothesis: the physics-based model can accurately predict long-term creep rupture in 316H, even when only provided with a test database for times as short as 5,000 hours. Moreover, the model performs significantly better when extrapolating than the conventional, empirical approach using Larson-Miller extrapolation. Widespread adoption of this new approach for extrapolating time-dependent material properties could accelerate the qualification of new materials by providing a much greater extrapolation factor, i.e. the ratio between the qualified material life and the longest required test. Even if this method does not completely replace the more conventional time-temperature parameter correlation it could be used to support staggered qualification programs by providing designers accurate, long-term predictions of key material properties for the component design based only on limited, short-term test data.

## 5.2 Future work

Our original plan was to use the Hu and Cocks model to generate the simulation database. To save time we proceeded in parallel implementing the new Hu and Cocks model and generating the simulation database using an existing model. While the results here demonstrate that the simple model was sufficiently accurate, the Hu and Cocks model better connects to observable microstructural parameters like the material chemistry, carbide distribution, etc. Repeating the process developed here using the Hu and Cocks model would yield a model that is easier to connect to characterization data via the parameter prior distributions. In turn, these informed priors would produce more accurate posterior parameter distributions and, potentially, could further improve the extrapolation results.

The general strategy developed here could be applied to other time dependent material properties like thermal aging, creep-fatigue, and even irradiation and environmental damage provided physics-based models for these behaviors can be developed. Thermal aging would be a very easy next target, as the Hu and Cocks model, or similar models for other materials, already predicts thermal aging along with creep.

Additional work could improve the repeatability of the process, mostly by examining the hyperparameters controlling the surrogate model training and inference processes. While the repeat variation in the example detailed in this report is acceptable, for material qualification codes and standards bodies may prefer a more repeatable process. However, the model calibration process will never be fully deterministic as the inference process will always remain stochastic.

In the longer term, relying on physics-based models means needing to develop a model for each substantially different material with significantly different deformation and damage mechanisms. While some model reuse is possible, for example between related materials of similar types, each substantially new material will require a new model. The time and cost of the research program to develop a new physics-based model, particularly one calibrated as described here, will not limit rapid material qualification as much as a long-term test program, but it does represent a cost that could potentially be reduced. Integrating the Bayesian inference, surrogate modeling, and CPFEE simulations into a single step using physics-informed machine learning might simplify the process of developing a new physics-based model. The idea would be to develop a “library” of mechanism models relating measurable microstructural parameters to creep deformation and failure. This library would include mechanisms relevant to a wide variety of material systems, not just the single material under consideration. These mechanism models could be combined with either full field CPFEE simulations or homogenized calculations to predict creep rupture or other long-term material behavior. Machine learning techniques could then be used to both down-select material mechanisms to only those that well-explain the creep rupture data and infer the resulting parameter distributions all in one step. Developing a modeling framework of this type could greatly reduce the time to develop a new physics-based model by leaving decisions on which mechanisms to include to the machine. However, it will require advances both in basic computational modeling to provide fast simulation results and in machine learning techniques to regularize the models in order to achieve.

## **Acknowledgements**

The research was sponsored by the U.S. Department of Energy (DOE), under Contract No. DE-AC02-06CH11357 with Argonne National Laboratory, managed and operated by UChicago Argonne LLC. Programmatic direction and funding was provided by the Nuclear Materials Discovery and Qualification Initiative (NMDQi) through the Department of Energy, Office of Nuclear Energy. The authors acknowledge the technical insight and direction provided by Michael D. McMurtrey and R. Allen Roach of Idaho National Laboratory as well as the oversight provided by Willettia Amos and Tansel Selekler at the DOE.

We gratefully acknowledge use of the Bebop cluster in the Laboratory Computing Resource Center at Argonne National Laboratory



## Bibliography

- [1] Jeffery A Aguiar, Andrea M Jokisaari, Matthew Kerr, and R Allen Roach. Bringing nuclear materials discovery and qualification into the 21 st century. *Nature Communications*, 11(1):1–3, 2020.
- [2] Robert Odette and Steven Zinkle. *Structural alloys for nuclear energy applications*. Elsevier, 2019.
- [3] F. R. Larson and J. Miller. A time-temperature relationship for rupture and creep stresses. *Transactions of the ASME*, 74:765–771, 1952.
- [4] S So Manson and Angela M Haferd. A linear time-temperature relation for extrapolation of creep and stress-rupture data. 1953.
- [5] B Wilshire, PJ Scharning, and R Hurst. A new approach to creep data assessment. *Materials Science and Engineering: A*, 510:3–6, 2009.
- [6] Daniel R Eno, George A Young, and T-L Sham. A unified view of engineering creep parameters. In *ASME Pressure Vessels and Piping Conference*, volume 48296, pages 777–792, 2008.
- [7] Kazuhiro Kimura, Hideaki Kushima, Kota Sawada, and Yoshiaki Toda. Region splitting analysis on creep strength enhanced ferritic steels. In *ASME Pressure Vessels and Piping Conference*, volume 42878, pages 647–654, 2007.
- [8] Omar Nassif, Timothy J Truster, Ran Ma, Kristine B Cochran, David M Parks, Mark Christian Messner, and TL Sham. Combined crystal plasticity and grain boundary modeling of creep in ferritic-martensitic steels: I. theory and implementation. *Modelling and Simulation in Materials Science and Engineering*, 27(7):075009, 2019.
- [9] T-L Sham and A Needleman. Effects of triaxial stressing on creep cavitation of grain boundaries. *Acta Metallurgica*, 31(6):919–926, jun 1983. ISSN 00016160. doi: 10.1016/0001-6160(83)90120-7.
- [10] A. Needleman and J. R. Rice. Plastic Creep Flow Effects in the Diffusive Cavitation of Grain Boundaries. *Acta Metallurgica*, 28(10):1315–1332, 1980. ISSN 00016160. doi: 10.1016/0001-6160(80)90001-2.
- [11] Erik Van Der Giessen and Viggo Tvergaard. Micromechanics of intergranular creep failure under cyclic loading. *Acta materialia*, 44(7):2697–2710, 1996.
- [12] Erik Van Der Giessen, M W D der Burg, A Needleman, and V Tvergaard. Void growth due to creep and grain boundary diffusion at high triaxialities. *Journal of the Mechanics and Physics of Solids*, 43(1):123–165, 1995. ISSN 00225096. doi: 10.1016/0022-5096(94)00059-E.
- [13] Patrick Onck and Erik Van Der Giessen. Growth of an initially sharp crack by grain boundary cavitation. *Journal of the Mechanics and Physics of Solids*, 47(1):99–139, 1998. ISSN 00225096. doi: 10.1016/S0022-5096(98)00078-7.

- [14] J. N. Hu and A. C.F. Cocks. Effect of creep on the Bauschinger effect in a polycrystalline austenitic stainless steel. *Scripta Materialia*, 128:100–104, 2017. ISSN 13596462. doi: 10.1016/j.scriptamat.2016.10.005.
- [15] Nailong Zhao, Anish Roy, Weizhe Wang, Liguang Zhao, and Vadim V. Silberschmidt. Coupling crystal plasticity and continuum damage mechanics for creep assessment in Cr-based power-plant steel. *Mechanics of Materials*, 130(December 2018):29–38, 2019. ISSN 01676636. doi: 10.1016/j.mechmat.2019.01.006. URL <https://doi.org/10.1016/j.mechmat.2019.01.006>.
- [16] Markian P. Petkov, Jianan Hu, Edmund Tarleton, and Alan C.F. Cocks. Comparison of self-consistent and crystal plasticity FE approaches for modelling the high-temperature deformation of 316H austenitic stainless steel. *International Journal of Solids and Structures*, 171:54–80, 2019. ISSN 00207683. doi: 10.1016/j.ijsolstr.2019.05.006.
- [17] Mark Christian Messner, Omar Nassif, Ran Ma, Timothy J Truster, Kristine Cochran, David Parks, and TL Sham. Combined crystal plasticity and grain boundary modeling of creep in ferritic-martensitic steels: II. the effect of stress and temperature on engineering and microstructural properties. *Modelling and Simulation in Materials Science and Engineering*, 27(7):075010, 2019.
- [18] C. M. Bishop and M. E. Tipping. Bayesian Regression and Classification. In J. A. K. Suykens, I. Horvath, S. Basu, and C. Micchelli, editors, *Advances in Learning Theory: Methods, Models and Applications*, pages 267–285. IOS Press, 2003. doi: 10.1017/S0030605311001384.
- [19] Rajesh Ranganath, Sean Gerrish, and David M. Blei. Black box variational inference. *Journal of Machine Learning Research*, 33:814–822, 2014. ISSN 15337928.
- [20] Mainak Sengupta and James E. Nestell. Correct and Extend Allowable Stress Values for 304 and 316 Stainless Steel. Technical report, ASME Standards Technology, LLC STP-NU-063, New York, 2013.
- [21] Jianan Hu, Graham Green, Simon Hogg, Rebecca Higginson, and Alan Cocks. Effect of microstructure evolution on the creep properties of a polycrystalline 316H austenitic stainless steel. *Materials Science and Engineering A*, 772(August 2019):138787, 2020. ISSN 09215093. doi: 10.1016/j.msea.2019.138787. URL <https://doi.org/10.1016/j.msea.2019.138787>.
- [22] Jianan Hu. *A theoretical study of creep deformation mechanisms of Type 316H stainless steel at elevated temperatures*. Phd dissertation, University of Oxford, 2015.
- [23] Robert J Asaro. Crystal plasticity. *Journal of Applied Mechanics*, 50:921–934, 1983.
- [24] Mark Messner, Armand Beaudoin, and Robert Dodds. Consistent crystal plasticity kinematics and linearization for the implicit finite element method. *Engineering Computations*, 2015.

- [25] A. Rovinelli, M. C. Messner, and T.-L. Sham. Initial microstructural model for creep-fatigue damage in grade 91 steel. Technical Report ANL-ART-202, Argonne National Laboratory, 2020.
- [26] Andrea Rovinelli, Mark C. Messner, David Parks, and T.-L. Sham. Accurate effective stress measures: Predicting creep life for 3D stresses using 2D and 1D creep rupture simulations and data. *Integrating Materials and Manufacturing Innovation*, In Press, 2021.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, pages 8024—8035, 2019.
- [28] Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method of Stochastic Optimization. In *International Conference on Learning Representations 2015*, pages 1–15, 2015.
- [29] Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine Learning Research*, 2018.
- [30] John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. *arXiv preprint arXiv:1506.05254*, 2015.









## **Applied Materials Division**

Argonne National Laboratory  
9700 South Cass Avenue, Bldg. 212  
Argonne, IL 60439

[www.anl.gov](http://www.anl.gov)



Argonne National Laboratory is a U.S. Department of Energy  
laboratory managed by UChicago Argonne, LLC